IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of:	**§ Group Art Unit: 2186**
Hofstee, et. al.	**§ Confirmation No.: 9220**
	§ Examiner: Iwashko, Lev
Serial No.: 10/697,897	**§**
	§ Attorney Docket No.
Filed: October 30, 2003	**§ AUS920030402US1**
	§
Title: System and Method for Sharing Memory by Heterogeneous Processors	**§ IBM Corporation**
	§ Intellectual Property Law
	§ Dept.
	§ 11400 Burnet Road
	§ Austin, Texas 78758

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Certificate of Mailing or Transmission

I hereby certify that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail in an envelope addressed to: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450 or electronically transmitted to the U.S. Patent and Trademark Office on the date shown below.

/Joseph T. Van Leeuwen, Reg. No. 44,383/ 10/10/2006
Joseph T. Van Leeuwen, Reg. No. 44,383 Date

REQUEST FOR CONTINUED EXAMINATION (RCE)

Sir:

A. INTRODUCTORY COMMENTS

In response to the Final Office Action having a mailing date of July 7, 2006, with a three-month shortened statutory period for response set to expire on October 10, 2006, Applicants respectfully request continued examination of the present application under 37 C.F.R. § 1.114, and request reconsideration of the outstanding rejections in light of the following remarks and amendments.

No extension of time is believed to be necessary. If, however, an extension of time is required, the extension is requested, and the undersigned hereby authorizes the Commissioner to charge any fees for this extension to IBM Corporation Deposit Account No. 09-0447.

B. AMENDMENTS TO THE CLAIMS

1. (Currently Amended) A memory shared by a plurality of heterogeneous processors, comprising:

the shared memory;

wherein the shared memory is accessible by one or more first processors that are adapted to process a first instruction set;

wherein the shared memory is accessible by one or more second processors that are adapted to process a second instruction set; and

a memory map corresponding to the shared memory, wherein the memory map includes cross-references between virtual addresses and real addresses, the memory map and the cross-references shared between the first processors and the second processors.

~~wherein the shared memory, the first processors, and the second processors are included on one silicon substrate and are connected using an on chip coherent multi-processor bus~~.

2. (Canceled)

3. (Currently Amended) The shared memory as described in claim [[2]] 1 further comprising:

an operating system that operates on one of the first processors, the first processor controlling the memory map.

4. (Original) The shared memory as described in claim 1 wherein each second processor further comprises:

a synergistic processing unit;

a local storage; and

a memory management unit, the memory management unit including a direct memory access controller.

5. (Original) The shared memory as described in claim 4 wherein at least one of the second processors use the direct memory access controller to access the shared memory.

6. (Original) The shared memory as described in claim 4 wherein the local storage is divided into a private storage and a non-private storage.

7. (Original) The shared memory as described in claim 6 wherein the non-private storage is included in the shared memory.

8. (Currently Amended) The shared memory as described [[on]] in claim [[2]] 1 wherein the memory map includes a plurality of regions, wherein at least one of the regions is selected from the group consisting of an external system memory region, a local storage aliases region, a TLB region, an MFC region, an operating system region, and an I/O devices region.

9. (Previously Canceled)

10. (Previously Canceled)

11. (Currently Amended) A method for sharing a memory between a plurality of heterogeneous processors, said method comprising:

receiving a memory request;

allocating a first memory partition on the shared memory that corresponds to the memory request, the first memory partition accessible by one or more first processors that are adapted to process a first instruction set;

assigning a second memory partition on the shared memory to one or more second processors that are adapted to process a second instruction set, wherein the first processors and the second processors are heterogeneous;

managing the first memory partition and the second memory partition using a common memory map; ~~and~~

wherein the common memory map includes a plurality of regions, wherein at least one of the regions is selected from the group consisting of an external system memory region, a local storage aliases region, a TLB region, an MFC region, an operating system region, and an I/O devices region; and

wherein the TLB region includes cross-references between virtual addresses and real addresses, the common memory map and the cross-references shared between the first processors and the second processors.

12. (Previously Canceled)

13. (Previously Amended) The method as described in claim 11 wherein one of the first processors includes an operating system whereby the first processor controls the common memory map.

14. (Canceled)

15. (Original) The method as described in claim 11 wherein at least one of the first processors is a Power PC and wherein at least one of the second processors is included in a synergistic processing unit.

16. (Original) The method as described in claim 15 wherein the shared memory corresponds to the synergistic processing unit.

17. (Original) The method as described in claim 11 wherein at least one of the second processors uses a direct memory access controller for accessing the shared memory.

18. (Currently Amended) A computer program product stored on a computer operable media for sharing a memory between a plurality of heterogeneous processors, said computer program product comprising:

means for allocating a first memory partition on the shared memory that corresponds to the memory request, the first memory partition accessible by one or more first processors that are adapted to process a first instruction set;

means for assigning a second memory partition on the shared memory to one or more second processors that are adapted to process a second instruction set, wherein the first processors and the second processors are heterogeneous; ~~and~~

means for managing the first memory partition and the second memory partition using a common memory map that includes a plurality of regions, wherein at least one of the regions is selected from the group consisting of an external system memory region, a local storage aliases

region, a TLB region, an MFC region, an operating system region, and an I/O devices region; and

wherein the TLB region includes cross-references between virtual addresses and real addresses, the common memory map and the cross-references shared between the first processors and the second processors.

19. (Previously Canceled)

20. (Previously Amended) The computer program product as described in claim 18 wherein one of the first processors includes an operating system whereby the first processor controls the common memory map.

21. (Previously Canceled)

22. (Original) The computer program product as described in claim 18 wherein at least one of the first processors is a Power PC and wherein at least one of the second processors is included in a synergistic processing unit.

23. (Original) The computer program product as described in claim 22 wherein the shared memory corresponds to the synergistic processing unit.

24. (Original) The computer program product as described in claim 18 wherein at least one of the second processors uses a direct memory access controller for accessing the shared memory.

25. (Currently Amended) A memory shared by a plurality of heterogeneous processors, comprising:

the memory, wherein the memory includes one or more non-private storage areas, the non-private storage areas included in one or more second processors that are adapted to process a second instruction set and access the memory; ~~and~~

wherein the shared memory is accessible by one or more first processors that are adapted to process a first instruction set and access the memory; and

a memory map corresponding to the shared memory, wherein the memory map includes cross-references between virtual addresses and real addresses, the memory map and the cross-references shared between the first processors and the second processors.

26. (Original) The shared memory as described in claim 25 wherein each second processor further comprises:

synergistic processing logic which uses private storage, the private storage not included in the shared memory; and

memory management logic for directly accessing the shared memory.

27. (Original) The shared memory as described in claim 25 further comprising:

memory mapping logic that corresponds to the shared memory, wherein the memory mapping logic is shared between the first processors and the second processors.

28. (Original) The shared memory as described in claim 27 further comprising:

an operating system that operates on one of the first processors, the first processor controlling the memory mapping logic.

29. (Original) The shared memory as described in claim 25 wherein one of the first processors configures each of the non-private storage areas.

30. (Newly Added) The shared memory as described in claim 1 further comprising:

wherein the shared memory, the first processors, and the second processors are included on one silicon substrate and are connected using an on chip coherent multi-processor bus.

C. REMARKS

Status of the Claims

Claims 1, 3-8, 11, 13, 15-18, 20, and 22-30 are currently present in the Application, and claims 1, 11, 18, and 25 are independent claims. Claims 1, 3, 8, 11, 18, and 25 have been amended, claim 2 has been cancelled, and claim 30 has been added.

Drawings

Applicants note with appreciation the Examiner's acceptance of Applicants' formal drawings filed concurrently with the application.

Claim Rejections

Claims 1-5 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over McCrory (International Publication WO 98/19238, hereinafter "McCrory") in view of Kumaki et al. (U.S. Patent No. 5,481,726, hereinafter "Kumaki") and Foster et al. (U.S. Patent No. 5,410,654, hereinafter "Foster"). Applicants respectfully traverse these rejections. Applicants have canceled dependent claim 2 in this response and, therefore, rejections to this claim are moot.

Applicants amended independent claim 1 to include the limitations found in original claim 2, as well as adding a limitation based upon Applicants' heterogeneous processors sharing the same *architectural address translation mechanism*. As amended, claim 1 is directed towards a memory shared by a plurality of heterogeneous processors with limitations comprising:

- the shared memory;

- wherein the shared memory is accessible by one or more first processors that are adapted to process a first instruction set;
- wherein the shared memory is accessible by one or more second processors that are adapted to process a second instruction set; and
- a memory map corresponding to the shared memory, wherein the memory map **includes cross-references between virtual addresses and real addresses** and is shared between the first processors and the second processors.

Amended claim 1 includes a limitation of a memory map that corresponds to the shared memory, "*wherein the memory map **includes cross-references between virtual addresses and real addresses**.*" Applicants' first and second processors use the same architectural address translation mechanism, and thus use the same cross-references between virtual addresses and real addresses. Applicants' original specification supports Applicants' amendment on page 47, lines 1-23 and, therefore, no new matter is added with such amendment.

The Office Action uses McCrory to reject Applicants' "memory map" limitation found in original claim 2 by suggesting that McCrory's common memory is the same as Applicants' memory map. After further review, however, McCrory's teachings are different than Applicants' invention in that McCrory uses the same ***physical*** memory addresses for both processor types, but does not use the same ***address translation cross-references*** as claimed by Applicants. Specifically, McCrory states:

> "In heterogeneous systems, however, in order to employ common data buses, address buses, and control buses to communicate between the various processors, **memory** and I/O, **some interfacing mechanism must be provided** between each family of processors and the common data, address, and control buses… Bus interface converter devices 322-326 provide physical and logical external conversions for coupling the processors within a

> processor family to implementation specific communication mechanism 328. **Each bus interface converter is unique to a processor family and possibly unique to a processor within a processor family.** The Intel 8046 and the Intel Pentium, for example, are in the same processor family but have different external interfaces and would probably require different bus interface converters." (page 12, lines 12-36, emphasis added)

As can be seen from the above excerpt, McCrory uses unique bus interface converters for each processor, and in no way teaches or suggests "*wherein the memory map includes cross-references between virtual addresses and real addresses, the memory map and the cross-references shared between the first processors and the second processors*" as claimed by Applicants. The Office Action does not suggest that Kumaki or Foster teach Applicants' memory map limitation, and indeed they do not.

Therefore, since McCrory, Kumaki, or Foster do not teach or suggest, either alone or in combination with each other, all the limitations of Applicants' claim 1 as amended, amended claim 1 is allowable over McCrory in view of Kumaki and further in view of Foster. Each of claims 3-5 depends, either directly or indirectly, upon allowable claim 1. Therefore, claims 3-5 are allowable for at least the same reasons as claim 1 is allowable.

Claims 6-7 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over McCrory in view of Parrish et al. (U.S. Patent No. 5,117,359) and further in view of Foster. Applicants respectfully traverse these rejections.

Claims 6-7 depend upon allowable claim 1 discussed above. The Office Action does not suggest that Parrish teaches or suggests the limitations included in Applicants' claim 1, and indeed Parrish does not teach such limitations. Therefore, claims 6-7 are allowable over McCrory in view of Parish and

Foster for at least the same reasons that claim 1 is allowable over McCrory in view of Kumaki and Foster.

Claim 8 stands rejected under 35 U.S.C. § 103(a) as being unpatentable over McCrory in view of Brown (NPL Document "The Design of ARMphetamin 2," hereinafter "Brown"), Foster, and Kumaki. Applicants respectfully traverse these rejections.

Claim 8 depends upon allowable claim 1 discussed above. The Office Action does not suggest that Brown teaches or suggests the limitations included in Applicants' claim 1, and indeed Brown does not teach such limitations. Therefore, claim 8 is allowable over McCrory in view of Brown, Foster, and Kumaki for at least the same reasons that claim 1 is allowable over McCrory in view of Kumaki and Foster.

Claims 11, 13, 15-18, 20, and 22-24 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Parrish in view of McCrory and further in view of Brown. Applicants respectfully traverse these rejections.

Applicants have amended independent claim 11 to distinctly claim that Applicants' heterogeneous processors share cross-references between virtual addresses and real addresses that are included in a memory map. As amended, claim 11 is directed towards a method for sharing a memory between a plurality of heterogeneous processors with limitations comprising:

- receiving a memory request;
- allocating a first memory partition on the shared memory that corresponds to the memory request, the first memory partition accessible by one or more first processors that are adapted to process a first instruction set;
- assigning a second memory partition on the shared memory to one or more second processors that are adapted to process a second instruction set, **wherein**

the first processors and the second processors are heterogeneous;

- managing the first memory partition and the second memory partition using a common memory map;
- wherein the **common memory map** includes a plurality of regions, wherein at least one of the regions is selected from the group consisting of an external system memory region, a local storage aliases region, a TLB region, an MFC region, an operating system region, and an I/O devices region; and
- **wherein the TLB region includes cross-references between virtual addresses and real addresses, the common memory map and the cross-references shared between the first processors and the second processors.**

Similar to claim 1 discussed above, Applicants have amended claim 11 to include a limitation of "*wherein the TLB region includes cross-references between virtual addresses and real addresses, the common memory map and the cross-references shared between the first processors and the second processors.*"

The Office Action uses Brown to reject Applicants' "plurality of regions" limitation. Brown, however, does not teach or suggest that Brown's TLB region is shared between heterogeneous processors because Brown teaches a single processor system, and not a system that includes heterogeneous processors. Parrish discloses the use of different local address translation mechanisms for each processor. Particularly, Parrish states that "*The mechanism may comprise partitioning RAMs at **each functional unit** which respond to an input address and readout a stored translation address*" (col. 4, lines 58-60, emphasis added). As such, Parrish teaches that each processor uses a separate, local translation mechanism, and therefore, does not teach that the cross-references are **shared**

between the first processors and the second processors. As discussed above, McCrory also does not teach such limitation.

Therefore, since Parrish, McCrory, or Brown do not teach or suggest, either alone or in combination with each other, all the limitations included in Applicants' claim 11 as amended, amended claim 11 is allowable over Parrish in view of McCrory and further in view of Brown. Claim 18 is a computer program product claim including similar limitations of claim 11 and, therefore, is allowable for at least the same reasons that claim 11 is allowable.

Each of claims 13, 15-17, 20, and 22-24 depends, either directly or indirectly, upon allowable claim 11 or claim 18. Therefore, claims 13, 15-17, 20, and 22-24 are also allowable for at least the same reasons as their respective dependent claims are allowable.

Claims 25-29 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over Parrish in view of Foster, McCrory, and Brown. Applicants respectfully traverse these rejections.

Applicants have amended claim 25 to further claim that a memory map includes virtual address to real address cross-references that are shared between heterogeneous processors. As amended, claim 25 includes the limitations of :

- the memory, wherein the memory includes one or more non-private storage areas, the non-private storage areas included in one or more second processors that are adapted to process a second instruction set and access the memory;
- wherein the shared memory is accessible by one or more first processors that are adapted to process a first instruction set and access the memory; and

- a memory map corresponding to the shared memory, wherein the memory map includes cross-references between virtual addresses and real addresses, the memory map and the cross-references shared between the first processors and the second processors.

As discussed above, Parrish, Foster, McCrory or Brown do not teach or suggest "*wherein the memory mapping logic includes cross-references between virtual addresses and real addresses, the memory mapping logic and the cross-references shared between the first processors and the second processors*" as claimed by Applicants. Therefore, since Parrish, Foster, McCrory or Brown do not teach or suggest, either alone or in combination with each other, all of the limitations included in Applicants' claim 25 as amended, amended claim 25 is allowable over Parrish in view of Foster, McCrory, and Brown.

Each of claims 26-29 depends, either directly or indirectly, upon allowable claim 25. Therefore, claims 26-29 are allowable for at least the same reasons that claim 25 is allowable.

New Claim

Applicants have added claim 30 in this response, which is similar to a limitation previously included in independent claim 1. Claim 30 is dependent upon claim 1 and, therefore, is allowable for at least the same reasons as claim 1 is allowable.

Conclusion

As a result of the foregoing, it is asserted by Applicants that the remaining claims in the Application are in condition for allowance, and Applicants respectfully request an early allowance of such claims.

Applicants respectfully request that the Examiner contact the Applicants' attorney listed below if the Examiner believes that such a discussion would be helpful in resolving any remaining questions or issues related to this Application.

Respectfully submitted,

By /Joseph T. Van Leeuwen, Reg. No. 44,383/

Joseph T. Van Leeuwen, Reg. No. 44,383
Van Leeuwen & Van Leeuwen
Attorney for Applicants
Telephone: (512) 301-6738
Facsimile: (512) 301-6742